arXiv:1701.07776v1 [stat.ME] 26 Jan 2017

Dependent Mixtures of Geometric Weights Priors

Spyridon J. Hatjispyros [1], Christos Merkatas

Department of Mathematics, University of the Aegean
Karlovassi, Samos, GR-832 00, Greece.

Abstract

A new approach on the joint estimation of partially exchangeable observations is presented by constructing pairwise dependence between m random density functions, each of which is modeled as a mixture of geometric stick breaking processes. This approach is based on a new random central masses version of the Pairwise Dependent Dirichlet Process prior mixture model (PDDP) first introduced in Hatjispyros et al. (2011). The idea is to create pairwise dependence through random measures that are location-preserving-expectations of Dirichlet random measures. Our contention is that mixture modeling with Pairwise Dependent Geometric Stick Breaking Process (PDGSBP) priors is sufficient for prediction and estimation purposes; moreover the associated Gibbs sampler is much faster and easier to implement than its Dirichlet Process based counterpart. To this respect, we provide a-priori-synchronized comparison studies under sparse m-scalable synthetic and real data examples.

Keywords: Bayesian nonparametric inference; Mixture of Dirichlet process; Geometric stick breaking weights; Geometric Stick Breaking Mixtures; Dependent Process.

1. Introduction. In Bayesian nonparametrics, the use of nonparametric priors such as the Dirichlet process (Ferguson, 1973), is justified by the assumption that the observations are exchangeble. That is, the distribution of $(X_1,\ldots,X_n)$ coincides with the distribution of $(X_{\pi(1)},\ldots,X_{\pi(n)})$, for all $\pi\in S(n)$, where $S(n)$ is the group of permutations of $\{1,\ldots,n\}$. However, in real life applications data are often partially exchangeable. For example the data may consist of independent observations sampled from m populations, or may be sampled from an experiment conducted in m different geographical places. This means that the joint law is invariant under permutations within m subgroups of observations $(X_{j,1},\ldots,X_{j,n_j})$, $j=1,\ldots,m$, then for all $\pi_j\in S(n_j)$

$$((X_{1,1},\ldots,X_{1,n_1}),\ldots,(X_{m,1},\ldots,X_{m,n_m}))\stackrel{\mathrm{d}}{=}((X_{1,\pi_1(1)},\ldots,X_{1,\pi_1(n_1)}),\ldots,(X_{m,\pi_m(1)},\ldots,X_{m,\pi_m(n_m)})). \quad (1)$$

When the exchangeability assumption fails one needs to use non exchangeable priors. There has been substantial research interest after the seminal work of MacEachern (1999) in the construction of dependent stochastic processes and using them as priors on Bayesian nonparametric models. These processes are distributions over a collection of measures indexed by values in

[1]Corresponding author. Tel.:+30 22730 82326
E-mail address: schatz@aegean.gr

some covariate space, such that the marginal distribution is described by a known nonparametric prior. The key idea is to induce dependence between a collection of random probability measures $(\mathbb{P}_j)_{1\leq j\leq m}$, where each $\mathbb{P}_j$ comes from a Dirichlet process (DP) with concentration parameter $c > 0$ and base measure P_0. Such random probability measures typically are used in mixture models to generate random densities $f_j(x) = \int_\Theta K(x|\theta)\mathbb{P}_j(d\theta)$ (Lo, 1984).

There is a variety of ways that a DP can be extended to a dependent DP. Most of them use the stick-breaking representation (Sethuraman, 1994), that is

$$\mathbb{P} = \sum_{k=1}^{\infty} w_k \delta_{\theta_k},$$

where $(\theta_k)_{k\geq 1}$ are independent and identically distributed from P_0 and $(w_k)_{k\geq 1}$ a stick breaking process, so if all $(v_k)_{k\geq 1}$ are idependent and identically distributed from the beta$(1, c)$ then $w_1 = v_1$ and for $k > 1$, $w_k = v_k \prod_{l<k}(1 - v_l)$, and introduce dependence through the weights and–or the atoms. In this direction, the most common dependent process is the Hierarchical DP (Teh et al., 2005) which induces dependence between a collection of random probability measures by letting the atoms to be common. A classical example for the use of dependent DP's is the Bayesian nonparametric regression problem where a random probability measure $\mathbb{P}_z$ is constructed for each covariate z and is of the type

$$\mathbb{P}_z = \sum_{k=1}^{\infty} w_k(z) \delta_{\theta_k(z)},$$

where $(w_k(z), \theta_k(z))$ are collection of processes indexed by the $z-$space. Extensions of DP models to dependent DP models can be found in De Iorio et al. (2004), Griffin and Steel (2006), and Dunson and Park (2008). A survey in the construction of dependent nonparametric priors can be found in Foti and Williamson (2014).

Recently there has been growing interest for the use of simpler random probability measures than the DP which although simpler are yet competent for Bayesian nonparametric density estimation. The geometric stick breaking (GSB) random probability measure (Fuentes–García, et al. 2010) has been used for density estimation purposes and has been shown to provide an efficient alternative to DP based mixture models. Some recent papers extend this nonparametric prior to a dependent nonparametric prior. In the direction of covariate dependent processes, GSB processes have been seen to provide an adequate model to the traditional dependent DP. For example, for Bayesian regression, Fuéntes–Garcia et al. (2009) propose a covariate dependent process based on random probability measures drawn from a GSB process. Ména et al. (2011) used GSB random probability measures in order to construct a purely atomic continuous time measure-valued process useful for the analysis of time series data. In this case the covariate $z \geq 0$ denotes the time that each observation is (discretely) recored and conditionally on z_i each observation is assumed to be drawn from a time-dependent nonparametric mixture model based on GSB processes. However, to the best of our knowledge random probability measures

drawn from a GSB process for modelling related density functions when samples from each density function are available have not been developed in the literature.

In this work we are going to construct pairwise dependent random probability measures based on GSB processes. That is, we are going to model a finite collection of m random distribution functions $(\mathbb{P}_j)_{1\leq j\leq m}$, where each $\mathbb{P}_j$ is a GSB random probability measure, such that there is a unique common component for each pair $(\mathbb{P}_j, \mathbb{P}_{j'})$ with $j \neq j'$. We are going to use these measures in the context of GSB mixture models, generating a collection of m GSB pairwise dependent random densities $(f_j(x))_{1\leq j\leq m}$. That is we are going to construct for any finite $m > 1$

$$(f_1, \ldots, f_m),$$

where each f_j, marginally is a random density function

$$f_j(x) = \int_\Theta K(x|\theta)\mathbb{P}_j(d\theta),$$

thus generalizing the GSB priors to a multivariate prior setting for partially exchangeable observations.

In the problem considered here, these random density functions $(f_j)_{1\leq j\leq m}$ are thought to be related and we aim to share information between groups to improve estimation of each density especially for those densities f_j for which the corresponding sample size n_j is small. In this direction, the main references include the work of Müller et al. (2014), Bulla et al. (2009), Kolossiatis et al. (2013) and Griffin et al. (2013). All these models have been proposed for the modeling of an arbitrary but finite number of random distribution functions, via a common part and an index specific idiosyncratic part so that for $0 < p_j < 1$ we have

$$\mathbb{P}_j = p_j\mathbb{P}_0 + (1 - p_j)\mathbb{P}_j^*,$$

where $\mathbb{P}_0$ is the common component to all other distributions and $\{\mathbb{P}_j^* : j = 1, \ldots, m\}$ are the idiosyncratic parts to each $\mathbb{P}_j$, and $\mathbb{P}_0, \mathbb{P}_j^* \overset{\text{iid}}{\sim} \text{DP}(c, P_0)$. In Lijoi et al. (2014) a more complex normalized random probability measure based on the σ–stable process is used for modeling dependent mixtures. Although similar (all models coincide only for the $m = 2$ case), these models are different from our model which is based on pairwise dependence of a sequence of random measures (Hatjispyros et al. 2011, 2016).

We are going to provide evidence through numerical experiments, that dependent GSB mixture models provide an efficient alternative to pairwise dependent DP (PDDP) priors. Firstly, we will randomize the existing PDDP model of Hatjispyros et al. (2011, 2016), by imposing gamma priors on its concentration masses, and then we will conduct a-priori synchronized density estimation comparison studies between the randomized PDDP model (rPDDP) and the pairwise dependent GSB process (PDGSBP) model using synthetic and real data examples.

This paper is organized as follows. In section 2 we will demonstrate the construction of dependent random densities based on a set of general random discrete distributions, using a

dependent model suggested by Hatjispyros et al. (2011). We demonstrate how specific choices of latent random variables can recover the model of Hatjispyros et al. and the dependent GSB model introduced in this paper. These latent variables will form the basis of a Gibbs sampler for posterior inference, given in Section 3. In Section 4 we resort to simulation. We provide a comparison study between the dependent DP model of Hatjispyros et al. with our dependent GSB model. For the dependent DP mixture model, we are going to model each $\mathbb{P}_{jl}$ as a random probability measure drawn from $\mathrm{DP}(c_{jl}, P_0)$. That is, $\mathbb{P}_{jl}, 1 \le j \le l \le m$ are mutually independent Dirichlet processes.

2. Preliminaries.

We consider an infinite real valued process $\{X_{ji} : 1 \le j \le m, i \ge 1\}$ defined over a probability space $(\Omega, \mathcal{F}, \mathrm{P})$, that is partially exchangeable as in (1). Let $\mathcal{P}$ denote the set of probability measures over $\mathbb{R}$; de Finetti proved that there exists a probability distribution Π over $\mathcal{P}^m$, that satisfies

$$\begin{aligned}
&\mathrm{P}\{X_{ji} \in A_{ji} : 1 \le j \le m, 1 \le i \le n_j\} \\
&= \int_{\mathcal{P}^m} \mathrm{P}\{X_{ji} \in A_{ji} : 1 \le j \le m, 1 \le i \le n_j \,|\, \mathbb{Q}_1, \ldots, \mathbb{Q}_m\} \, \Pi(d\mathbb{Q}_1, \ldots, d\mathbb{Q}_m) \\
&= \int_{\mathcal{P}^m} \prod_{j=1}^{m} \mathrm{P}\{X_{ji} \in A_{ji} : 1 \le i \le n_j \,|\, \mathbb{Q}_j\} \, \Pi(d\mathbb{Q}_1, \ldots, d\mathbb{Q}_m) \\
&= \int_{\mathcal{P}^m} \prod_{j=1}^{m} \left\{ \prod_{i=1}^{n_j} \mathbb{Q}_j(A_{ji}) \right\} \Pi(d\mathbb{Q}_1, \ldots, d\mathbb{Q}_m) \,.
\end{aligned}$$

The de Finetti measure Π represents a prior distribution over partially exchangeable observations.

We start off by describing the model as it was in Hatjispyros et al. (2011) and then proceed to the specific details for the case of GSB random measures. So, we have the set of random density functions generated via

$$f_j(x) := f_j(x|\, \mathbb{G}_{jl}, 1 \le l \le m) \;=\; \sum_{l=1}^{m} p_{jl}\, g_{j\,l}(x|\, \mathbb{G}_{jl}), \quad 1 \le j \le m, \tag{2}$$

where $\sum_{l=1}^{m} p_{jl} = 1$ a.s., the random densities satisfy $g_{jl} = g_{lj}$ and are independent mixtures of GSB processes under the slightly altered definition

$$\mathbb{P}_{jl} = \sum_{k=1}^{\infty} q_{jlk} \delta_{\theta_{jlk}} \text{ with } q_{jlk} = \lambda_{jl}(1-\lambda_{jl})^{k-1},\ \lambda_{jl} \sim h(\,\cdot\,; \xi_{jl}),\ \theta_{jlk} \overset{\text{iid}}{\sim} P_0, \tag{3}$$

where h is a parametric density supported over the interval $(0,1)$ and depending on some parameter $\xi_{jl} \in \Xi$, and P_0 the central measure for which $\mathbb{E}(\mathbb{G}_{jl}(A)) = P_0(A)$ for all Borel sets A of $\mathbb{R}$. Then

$$g_{jl}(x) := g_{jl}(x|\, \mathbb{G}_{jl}) = \int_{\Theta} K(x|\, \theta)\, \mathbb{G}_{jl}(d\theta),$$

for some kernel density $K(\,\cdot\,|\,\cdot\,)$ and $\{\mathbb{G}_{jl} : 1 \le j, l \le m\}$ form a matrix $\mathbb{G}$ of random distributions with $\mathbb{G}_{jl} = \mathbb{G}_{lj}$ for $j > l$ and each other element is an independent GSB process. The

random densities $(f_j)_{1\leq j\leq m}$ are dependent mixtures of the dependent random measures

$$\mathbb{Q}_j \;=\; \sum_{l=1}^{m} p_{jl}\,\mathbb{G}_{jl}, \quad 1\leq j\leq m. \tag{4}$$

In matrix notation

$$\mathbb{Q} \;=\; (p\otimes\mathbb{G})\,\mathbf{1}, \tag{5}$$

where $p=(p_{jl})$ is the matrix of random selection weights, and $p\otimes\mathbb{G}$ is the Hadamard product of the two matrices defined as $(p\otimes\mathbb{G})_{jl}=p_{jl}\mathbb{G}_{jl}$. By letting $\mathbf{1}$ to denote the $m\times 1$ matrix of ones it is that the jth element of vector $\mathbb{Q}$ is given by equation (4).

Following a univariate construction of geometric slice sets (Fuentes–García et al. 2010), we define the stochastic variables $\mathbf{N}=(N_{ji})$ for $1\leq i\leq n_j$ and $1\leq j\leq m$, where N_{ji} is an almost surely finite random variable of mass f_N possibly depending on parameters, associated with the sequential slice set $\mathcal{S}_{ji}=\{1,\ldots,N_{ji}\}$. Following Hatjispyros et al. (2011, 2016) we introduce:

1. The GSB mixture selection variables $\boldsymbol{\delta}=(\delta_{ji})$; for an observation x_{ji} that comes from f_j, δ_{ji} selects the GSB mixture $g_{ji}(x)$ that the observation came from.

2. The clustering variables $\boldsymbol{d}=(d_{ji})$; for an observation x_{ji} that comes from f_j, given δ_{ji}, d_{ji} allocates the component of the GSB mixture $g_{j\delta_{ji}}(x)$ that x_{ji} came from.

Proposition 1. *Suppose that the clustering variables (d_{ji}) conditionally on the slice variables (N_{ji}) are having the discrete uniform distribution over the sets $(\mathcal{S}_{ji})$ that is $d_{ji}|N_{ji}\sim\mathcal{DU}(\mathcal{S}_{ji})$, then*

$$f_j(x_{ji},N_{ji}=r)=\frac{1}{r}\sum_{l=1}^{m}p_{jl}f_N(r;\lambda_{jl})\sum_{k=1}^{r}K(x_{ji}|\theta_{jlk}), \tag{6}$$

and

$$f_j(x_{ji},N_{ji}=r,d_{ji}=k|\delta_{ji}=l)=\frac{1}{r}f_N(r;\lambda_{jl})\,\mathcal{I}(k\leq r)\,K(x_{ji}|\theta_{jlk}). \tag{7}$$

Proof. Starting from the N_{ji}-augmented random densities we have

$$\begin{aligned}
f_j(x_{ji},N_{ji}=r) &= \sum_{l=1}^{m}f_j(x_{ji},N_{ji}=r,\delta_{ji}=l)=\sum_{l=1}^{m}p_{jl}\,f_j(x_{ji},N_{ji}=r|\delta_{ji}=l)\\
&= \sum_{l=1}^{m}p_{jl}\sum_{k=1}^{\infty}f_j(x_{ji},N_{ji}=r,d_{ji}=k|\delta_{ji}=l)\\
&= \sum_{l=1}^{m}p_{jl}f_j(N_{ji}=r|\delta_{ji}=l)\sum_{k=1}^{\infty}f_j(d_{ji}=k|N_{ji}=r)f_j(x_{ji}|d_{ji}=k,\delta_{ji}=l).
\end{aligned}$$

Because $f_j(N_{ji}=r|\delta_{ji}=l)=f_N(r;\lambda_{jl})$ and $f_j(x_{ji}|d_{ji}=k,\delta_{ji}=l)=K(x_{ji}|\theta_{jlk})$, the last equation gives

$$\begin{aligned}
f_j(x_{ji},N_{ji}=r) &= \sum_{l=1}^{m}p_{jl}f_N(r;\lambda_{jl})\sum_{k=1}^{\infty}\frac{1}{r}\mathcal{I}(k\leq r)K(x_{ji}|\theta_{jlk})\\
&= \frac{1}{r}\sum_{l=1}^{m}p_{jl}f_N(r;\lambda_{jl})\sum_{k=1}^{r}K(x_{ji}|\theta_{jlk}).
\end{aligned}$$

Augmenting further with the variables d_{ji} and δ_{ji} yields

$$f_j(x_{ji}, N_{ji} = r, d_{ji} = k, \delta_{ji} = l) = \frac{1}{r}\, p_{jl}\, f_N(r; \lambda_{jl})\, \mathcal{I}(k \leq r)\, K(x_{ji}|\theta_{jlk}).$$

Because $\mathrm{P}(\delta_{ji} = l) = p_{jl}$, the last equation leads to equation (7) and the proposition follows. □

The following proposition gives a multivariate analogue of equation (2) in Fuentes–García, et al. (2010):

Proposition 2. *Given the random set $\mathcal{S}_{ji}$, the random functions in (2) become finite mixtures of a.s. finite equally weighted mixtures of the $K(\,\cdot\,|\,\cdot\,)$ probability kernels, that is*

$$f_j(x_{ji}|N_{ji} = r) = \sum_{l=1}^{m} \mathcal{W}(r; \lambda_{jl}) \sum_{k=1}^{r} \frac{1}{r} K(x_{ji}|\theta_{jlk}), \tag{8}$$

with

$$\mathcal{W}(r; \lambda_{jl}) = \frac{p_{jl} f_N(r; \lambda_{jl})}{\sum_{l'=1}^{m} p_{jl'} f_N(r; \lambda_{jl'})}.$$

Proof. Marginalizing the joint of x_{ji} and N_{ji} with respect to x_{ji} we obtain

$$f_j(N_{ji} = r) = \sum_{l=1}^{m} p_{jl} f_N(r; \lambda_{jl}).$$

Then dividing equation (6) with the probability that N_{ji} equals r we obtain equation (8). □

We note that the one-dimensional model introduced in Fuentes–García, et al. (2010) under our notation has the representation

$$f_j(x_{ji}|N_{ji} = r, \delta_{ji} = l) = \sum_{k=1}^{r} \frac{1}{r} K(x_{ji}|\theta_{jlk}).$$

3. The Model.

Marginalizing (7) with respect to (N_{ji}, d_{ji}) we obtain

$$f_j(x_{ji}|\delta_{ji} = l) = \sum_{k=1}^{\infty} \left(\sum_{r=k}^{\infty} \frac{1}{r} f_N(r; \lambda_{jl}) \right) K(x_{ji}|\theta_{jlk}). \tag{9}$$

The quantity inside the parenthesis on the right-hand side of the previous equation is $f_j(d_{ji}|\delta_{ji} = l)$. Following Fuentes–García, et al. (2010), we substitute $f_N(r; \lambda_{jl})$ with the negative binomial distribution $\mathcal{NB}(r; 2, \lambda_{jl})$, i.e.

$$f_N(r; \lambda_{jl}) = r\lambda_{jl}^2 (1 - \lambda_{jl})^{r-1} \mathcal{I}(r \geq 1), \tag{10}$$

then equation (9) becomes

$$f_j(x_{ji}|\delta_{ji}=l)=\sum_{k=1}^{\infty}q_{jlk}K(x_{ji}|\theta_{jlk}) \text{ with } q_{jlk}=\lambda_{jl}(1-\lambda_{jl})^{k-1},$$

and the f_j random densities take the form of a finite mixture of GSB mixtures

$$f_j(x_{ji})=\sum_{l=1}^{m}p_{jl}\sum_{k=1}^{\infty}q_{jlk}K(x_{ji}|\theta_{jlk}).$$

We denote the set of observations along the m groups as $\boldsymbol{x}=(x_{ji})$ and with $\boldsymbol{x}_j$ the set of observations in the jth group. The three sets of latent variables in the jth group will be denoted as $\boldsymbol{N}_j$ for the slice variables, $\boldsymbol{d}_j$ for the clustering variables, and finally $\boldsymbol{\delta}_j$ for the set of GSB mixture allocation variables. From now on, we are going to leave the auxiliary variables unspecified; especially for δ_{ji} we use the notation

$$\delta_{ji}=(\delta_{ji}^1,\ldots,\delta_{ji}^m)\in\{\mathbf{e}_1,\ldots,\mathbf{e}_m\} \text{ with } \mathrm{P}(\delta_{ji}=\mathbf{e}_l)=p_{jl},$$

where $\mathbf{e}_l$ denotes the usual basis vector having its only nonzero component equal to 1 at position l. Hence, for a sample of size n_1 from f_1, a sample of size n_2 from f_2, etc., a sample of size n_m from f_m we can write the full likelihood as a multiple product:

$$\begin{aligned}f(\boldsymbol{x},\boldsymbol{N},\boldsymbol{d}\,|\,\boldsymbol{\delta}) &= \prod_{j=1}^{m}f(\boldsymbol{x}_j,\boldsymbol{N}_j,\boldsymbol{d}_j\,|\,\boldsymbol{\delta}_j)\\ &= \prod_{j=1}^{m}\prod_{i=1}^{n_j}\mathcal{I}(d_{ji}\le N_{ji})\prod_{l=1}^{m}\left\{\lambda_{jl}^2(1-\lambda_{jl})^{N_{ji}-1}K(x_{ji}|\,\theta_{jld_{ji}})\right\}^{\delta_{ji}^l}.\end{aligned}$$

In a hierarchical fashion, using the auxiliary variables, we have for $j=1,\ldots,m$ and $i=1,\ldots,n_j$,

$$\begin{aligned}&x_{ji},N_{ji}\,|\,d_{ji},\delta_{ji},(\theta_{jr\delta_{ji}})_{1\le r\le m},\lambda_{j\delta_{ji}}\overset{\text{ind}}{\sim}\prod_{r=1}^{m}\left\{\lambda_{jl}^2(1-\lambda_{jl})^{N_{ji}-1}K(x_{ji}|\theta_{jrd_{ji}})\right\}^{\delta_{ji}^r}\mathcal{I}(N_{ji}\ge d_{ji})\\ &d_{ji}\,|\,N_{ji}\overset{\text{ind}}{\sim}\mathcal{DU}(\mathcal{S}_{ji}),\ \ \mathrm{P}(\delta_{ji}=\mathbf{e}_l)=p_{jl}\\ &q_{jik}=\lambda_{ji}(1-\lambda_{ji})^{k-1},\ \ \theta_{jik}\overset{\text{iid}}{\sim}P_0,\ \ k\in\mathbb{N},\end{aligned}$$

3.1 The PDGSBP Covariance.

In this section we find the covariance between $f_j(x)$ and $f_i(x)$. First we provide the following lemma.

Lemma 1. *Let $g_{\mathbb{G}}(x)=\int_\Theta K(x|\theta)\mathbb{G}(d\theta)$ be a random density, with $\mathbb{G}=\lambda\sum_{j=1}^{\infty}(1-\lambda)^{j-1}\delta_{\theta_j}$ and $\theta_j\overset{\text{iid}}{\sim}P_0$, then*

$$\mathbb{E}[g_{\mathbb{G}}(x)^2]=\left(\frac{1}{2-\lambda}\right)\left\{\lambda\int_\Theta K(x|\theta)^2P_0(d\theta)+2(1-\lambda)\left(\int_\Theta K(x|\theta)P_0(d\theta)\right)^2\right\}.$$

Proof. Because $g_{\mathbb{G}}(x) = \lambda \sum_{j=1}^{\infty}(1-\lambda)^{j-1}K(x|\theta_j)$, we have

$$\begin{aligned}
\mathbb{E}\left\{g_{\mathbb{G}}(x)^2\right\} &= \lambda^2\,\mathbb{E}\left\{\left(\sum_{j=1}^{\infty}(1-\lambda)^{j-1}K(x|\theta_j)\right)^2\right\} \\
&= \lambda^2\left\{\sum_{j=1}^{\infty}(1-\lambda)^{2j-2}\,\mathbb{E}\left[K(x|\theta_j)^2\right] + 2\sum_{k=2}^{\infty}\sum_{j=1}^{k-1}(1-\lambda)^{j+k-2}\,\mathbb{E}[K(x|\theta_j)K(x|\theta_k)]\right\} \\
&= \lambda^2\left\{\sum_{j=1}^{\infty}(1-\lambda)^{2j-2}\mathbb{E}\left[K(x|\theta)^2\right] + 2\sum_{k=2}^{\infty}\sum_{j=1}^{k-1}(1-\lambda)^{j+k-2}\mathbb{E}[K(x|\theta)]^2\right\} \\
&= \lambda^2\left\{\frac{1}{\lambda(2-\lambda)}\mathbb{E}\left[K(x|\theta)^2\right] + 2\,\frac{1-\lambda}{\lambda^2(2-\lambda)}\mathbb{E}[K(x|\theta)]^2\right\},
\end{aligned}$$

which gives the desired result. □

Proposition 3. *It is that*

$$\mathrm{Cov}(f_j(x), f_i(x)) \;=\; p_{ji}\,p_{ij}\mathrm{Var}\left(\int_{\Theta}K(x|\theta)\mathbb{G}_{ji}(d\theta)\right),$$

with

$$\mathrm{Var}\left(\int_{\Theta}K(x|\theta)\mathbb{G}_{ji}(d\theta)\right) = \frac{\lambda_{ji}}{2-\lambda_{ji}}\mathrm{Var}(K(x|\theta)).$$

Proof. The random densities $f_i(x) = \sum_{l=1}^{m} p_{il}\,g_{il}(x)$ and $f_j(x) = \sum_{l=1}^{m} p_{jl}\,g_{jl}(x)$ depend to each other through the random measure $\mathbb{G}_{ji}$, therefore

$$\mathbb{E}[f_i(x)f_j(x)] = \mathbb{E}[\,\mathbb{E}(f_i(x)f_j(x)|\mathbb{G}_{ji})\,] = \mathbb{E}\{\,\mathbb{E}[f_i(x)|\mathbb{G}_{ji}]\,\mathbb{E}[f_j(x)|\mathbb{G}_{ji}]\,\}, \tag{11}$$

and

$$\begin{aligned}
\mathbb{E}[f_j(x)|\mathbb{G}_{ji}] &= \sum_{l\neq i}p_{jl}\,\mathbb{E}[g_{jl}(x)] + p_{ji}g_{ji}(x) = (1-p_{ji})\,\mathbb{E}[K(x|\theta)] + p_{ji}g_{ji}(x) \\
\mathbb{E}[f_i(x)|\mathbb{G}_{ji}] &= \sum_{l\neq j}p_{il}\,\mathbb{E}[g_{il}(x)] + p_{ij}g_{ji}(x) = (1-p_{ij})\,\mathbb{E}[K(x|\theta)] + p_{ij}g_{ji}(x)\,.
\end{aligned}$$

Substituting back to equation (11) one obtains

$$\mathbb{E}[f_i(x)f_j(x)] = (1-p_{ij}p_{ji})\,\mathbb{E}[K(x|\theta)]^2 + p_{ij}p_{ji}\,\mathbb{E}\left[g_{ji}(x)^2\right].$$

Using lemma 1, the last equation becomes

$$\mathbb{E}[f_i(x)f_j(x)] = \frac{\lambda_{ji}p_{ji}p_{ij}}{2-\lambda_{ji}}\left\{\mathbb{E}[K(x|\theta)^2] - \mathbb{E}[K(x|\theta)]^2\right\} + \mathbb{E}[K(x|\theta)]^2,$$

or that

$$\mathrm{Cov}(f_j(x), f_i(x)) \;=\; \frac{\lambda_{ji}p_{ji}\,p_{ij}}{2-\lambda_{ji}}\mathrm{Var}(K(x|\theta)).$$

The desired result, comes from the fact that

$$\begin{aligned}\mathrm{Var}\left(\int_{\Theta} K(x|\theta)\mathbb{G}_{ji}(d\theta)\right) &= \left\{\frac{\lambda_{ji}}{2-\lambda_{ji}}\mathbb{E}[K(x|\theta)^2] + \frac{2(1-\lambda_{ji})}{2-\lambda_{ji}}\mathbb{E}[K(x|\theta)]^2\right\} - \mathbb{E}[K(x|\theta)]^2 \\ &= \frac{\lambda_{ji}}{2-\lambda_{ji}}\left(\mathbb{E}[K(x|\theta)^2] - \mathbb{E}[K(x|\theta)]^2\right).\end{aligned}$$

□

Suppose now that $(f_j^{\mathcal{D}}(x))_{1\leq j\leq m}$ and $(f_j^{\mathcal{G}}(x))_{1\leq j\leq m}$ are two collections of m DP and GSB pairwise dependent random densities respectively, i.e. $f_j^{\mathcal{D}}(x) = \sum_{l=1}^m p_{jl} g_{jl}(x|\mathbb{P}_{jl})$ and $f_j^{\mathcal{G}}(x) = \sum_{l=1}^m p_{jl} g_{jl}(x|\mathbb{G}_{jl})$. In Hatjispyros et al. 2011 it is shown that

$$\mathrm{Cov}(f_j^{\mathcal{D}}(x), f_i^{\mathcal{D}}(x)) = \frac{p_{ij}p_{ji}}{1+c_{ji}}\mathrm{Var}(K(x|\theta)).$$

The latter equation combined with proposition 3 gives the following:

Proposition 4. *For $\lambda_{ji} = (1+c_{ji})^{-1}$ it is that*

$$\mathrm{Cov}(f_j^{\mathcal{G}}(x), f_i^{\mathcal{G}}(x)) < \mathrm{Cov}(f_j^{\mathcal{D}}(x), f_i^{\mathcal{D}}(x)).$$

4. The PDGSBP Gibbs Sampler.

In this section we are going to describe the PDGSBP Gibbs sampler for estimating the model. The details for the sampling algorithm of the PDDP model can be found in Hatjispyros et al. (2011, 2016). At each iteration we will sample variables,

$$\begin{aligned}&\theta_{jlk}, 1\leq j\leq l\leq m,\ 1\leq k\leq N^*,\\ &d_{ji}, N_{ji}, \delta_{ji}, 1\leq j\leq m,\ 1\leq i\leq n_j,\\ &p_{jl}, 1\leq j\leq m, 1\leq l\leq m,\end{aligned}$$

with $N^* = \max_{j,i} N_{ji}$ almost surely finite.

1. For the locations of the random measures for $k = 1,\ldots,d^*$ where $d^* = \max_{j,i} d_{ji}$, it is that

$$f(\theta_{jlk}|\cdots) \propto f(\theta_{jlk}) \begin{cases} \displaystyle\prod_{i=1}^{n_j} K(x_{ji}|\theta_{jlk})^{\mathcal{I}(\delta_{ji}=\mathbf{e}_l,\, d_{ji}=k)} \prod_{i=1}^{n_l} K(x_{li}|\theta_{jlk})^{\mathcal{I}(\delta_{li}=\mathbf{e}_j,\, d_{li}=k)} & l>j\,, \\ \displaystyle\prod_{i=1}^{n_j} K(x_{ji}|\theta_{jjk})^{\mathcal{I}(\delta_{ji}=\mathbf{e}_j,\, d_{ji}=k)} & l=j\,. \end{cases}$$

If $N^* > d^*$ we sample additional locations $\theta_{jl,d^*+1},\ldots,\theta_{jl,N^*}$ independently from the prior.

2. Here we sample the allocation variables d_{ji} and the mixture component indicator variables δ_{ji} as a block. For $j = 1,\ldots,m$ and $i = 1,\ldots,n_j$, we have

$$\mathrm{P}(d_{ji}=k, \delta_{ji}=\mathbf{e}_l\,|N_{ji}=r,\cdots) \;\propto\; p_{jl}\,K(x_{ji}|\theta_{jlk})\,\mathcal{I}(l\leq m)\,\mathcal{I}(k\leq r).$$

3. The slice variables N_{ji} have full conditional distributions given by

$$\mathrm{P}(N_{ji} = r \,|\, \delta_{ji} = \mathbf{e}_l, d_{ji} = l, \cdots) \propto (1-\lambda_{jl})^r \, \mathcal{I}(r \geq l),$$

which are truncated geometric distributions over the set $\{l, l+1, \ldots\}$.

4. The full conditional for $j = 1, \ldots, m$ for the selection probabilities $\boldsymbol{p}_j = (p_{j1}, \ldots, p_{jm})$, under a Dirichlet prior $f(\boldsymbol{p}_j \,|\, \boldsymbol{a}_j) \propto \prod_{l=1}^m p_{jl}^{a_{jl}-1}$, with hyperparameter $\boldsymbol{a}_j = (a_{j1}, \ldots, a_{jm})$, is Dirichlet

$$f(\boldsymbol{p}_j \,|\, \cdots) \;\propto\; \prod_{l=1}^{m} p_{jl}^{a_{jl}+\sum_{i=1}^{n_l} \mathcal{I}(\delta_{ji}=\mathbf{e}_l)-1}.$$

5. Here we update the geometric probabilities (λ_{jl}) of the GSB measures. For $1 \leq j \leq l \leq m$, it is that

$$f(\lambda_{jl}|\cdots) \propto f(\lambda_{jl}) \begin{cases} \displaystyle\prod_{i=1}^{n_j} \left\{\lambda_{jl}^2 (1-\lambda_{jl})^{N_{ji}-1}\right\}^{\mathcal{I}(\delta_{ji}=\mathbf{e}_l)} \prod_{i=1}^{n_l} \left\{\lambda_{jl}^2 (1-\lambda_{jl})^{N_{li}-1}\right\}^{\mathcal{I}(\delta_{li}=\mathbf{e}_j)} & l > j \\ \displaystyle\prod_{i=1}^{n_j} \left\{\lambda_{jj}^2 (1-\lambda_{jj})^{N_{ji}-1}\right\}^{\mathcal{I}(\delta_{ji}=\mathbf{e}_j)} & l = j\,. \end{cases}$$

To complete the model, we assign priors to the geometric probabilities. For a fair comparison between the two models, we apply $\lambda_{jl} = (1+c_{jl})^{-1}$ transformed priors. So, by placing gamma priors $c_{jl} \sim \mathcal{G}(a_{jl}, b_{jl})$ over the concentration masses c_{jl} of the PDDP model, we have

$$f(\lambda_{jl}) = \mathcal{TG}(\lambda_{jl} \,|\, a_{jl}, b_{jl}) \propto \lambda_{jl}^{-(a_{jl}+1)} e^{-b_{jl}/\lambda_{jl}} (1-\lambda_{jl})^{a_{jl}-1} \, \mathcal{I}(0 < \lambda_{jl} < 1).$$

In the appendix, we give the full conditionals for λ_{jl}'s, their corresponding embedded Gibbs sampling schemes, and the sampling algorithm for the concentration masses.

7. The Complexity of the rPDDP and PDGSBP samplers

The main difference between the two samplers in terms of execution time, comes from the blocked sampling of the clustering and the mixture indicator variables d_{ji} and δ_{ji}.

The rPDDP model: The state space of the variable (d_{ji}, δ_{ji}) conditionally on the slice variable u_{ji} is

$$(d_{ji}, \delta_{ji})(\Omega) = \cup_{l=1}^m \left(A_{w_{jl}}(u_{ji}) \times \{\mathbf{e}_l\}\right),$$

where $A_{w_{jl}}(u_{ji}) = \{r \in \mathbb{N} : u_{ji} < w_{jlr}\}$ is the a.s. finite slice set corresponding to the observation x_{ji}. At each iteration of the Gibbs sampler, we have $m(m+1)/2$ vectors of stick-breaking weights $\mathbf{w}_{jl}$, each of length N_{jl}^*; where $N_{jl}^* \sim 1 + \text{Poisson}(-c_{jl} \log u_{jl}^*)$ with c_{jl} being the concentration parameter of the Dirichlet process $\mathbb{P}_{jl}$ and u_{jl}^* being the minimum of the slice variables in densities f_j and f_l. Algorithm 1 gives the blocked sampling procedure of the clustering and mixture indicator variables. An illustration of the effect of the slice variable u_{ji} is given in figure 1(a).

Algorithm 1 : rPDDP

1: **procedure** SAMPLE (d_{ji}, δ_{ji})
2: **for** random densities f_j, $j = 1$ to m **do**
3: **for** each data point $x_{ji} \in f_j$ $i = 1$ to n_j **do**
4: **for** each mixture component $K(x_{ji}|\theta_{jl})$, $l = 1$ to m **do**
5: Construct slice sets $A_{w_{jl}}(u_{ji})$
6: **end for**
7: Sample $(d_{ji} = k, \delta_{ji} = r| \cdots) \propto K(x_{ji}|\theta_{jrk})\, \mathcal{I}\left((k,r) \in \cup_{l=1}^{m} \left(A_{w_{jl}}(u_{ji}) \times \{\mathbf{e}_l\}\right)\right)$
8: **end for**
9: **end for**
10: **end procedure**

Since the weights forming the stick-breaking representation are not in an ordered form, the construction of the slice sets in step 5 of Algorithm 1 requires a complete search in the array where the weights are stored. This operation is done in $\mathcal{O}(N_{jl}^*)$ time. For the sampling of the d_{ji} and δ_{ji} variables in step 6, the choice of their value is an element from the union $\cup_{l=1}^{m} \left(A_{w_{jl}}(u_{ji}) \times \{\mathbf{e}_l\}\right)$. This means that the rPDDP algorithm for each j, must create m slice sets which require M_{jl}^* comparisons each. The worst case scenario is that the sampled (d_{ji}, δ_{ji}) is the last element of $\cup_{l=1}^{m} \left(A_{w_{jl}}(u_{ji}) \times \{\mathbf{e}_l\}\right)$. Thus, the DP based procedure of sampling (d_{ji}, δ_{ji}) is of order

$$\mathcal{O}\left(m^2 n_j N_{jl}^* \sum_{l=1}^{m} |A_{w_{jl}}(u_{ji})|\right) = \mathcal{O}\left(N_{jl}^* \sum_{l=1}^{m} |A_{w_{jl}}(u_{ji})|\right).$$

The PDGSBP model: The state space of the variable (d_{ji}, δ_{ji}) conditionally on the slice variable N_{ji} is

$$(d_{ji}, \delta_{ji})(\Omega) = \cup_{l=1}^{m} \left(\mathcal{S}_{ji} \times \{\mathbf{e}_l\}\right).$$

In the GSB case, the slice variable has a different rolê. It indicates at which random point the search for the appropriate d_{ji} will stop. In Figure 1(b) we illustrate this argument. In Algorithm 2 the worst case scenario is that the sampled (d_{ji}, δ_{ji}) will be the last element of $\cup_{l=1}^{m} \left(\mathcal{S}_{ji} \times \{\mathbf{e}_l\}\right)$. Thus, the GSB based procedure of sampling (d_{ji}, δ_{ji}) is of order

$$\mathcal{O}\left(m^2 n_j N_{jl}\right) = \mathcal{O}\left(N_{jl}\right).$$

Algorithm 2 : PDGSBP

procedure SAMPLE (d_{ji}, δ_{ji})
  for random densities f_j, $j = 1$ to m **do**
    for each data point $x_{ji} \in f_j$ $i = 1$ to n_j **do**
      for each mixture component $K(x_{ji}|\theta_{jl})$, $l = 1$ to m **do**
        Sample $(d_{ji} = k, \delta_{ji} = r|\cdots) \propto K(x_{ji}|\theta_{jrk})\,\mathcal{I}(k \leq N_{ji})\,\mathcal{I}(r \leq m)$
      end for
    end for
  end for
end procedure

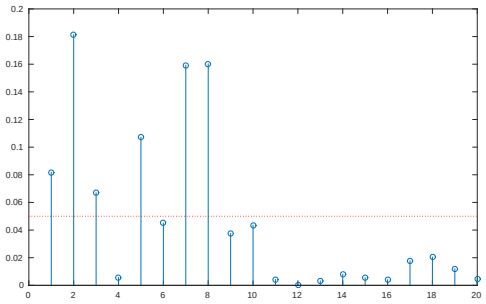

(a) Stick-breaking weights for some $N^*_{jl} = 20$. The red dashed line represents the slice variable $u_{ji} = 0.05$. The algorithm must check all the N^*_{jl} values to accept those that they satisfy $u_{ji} < w_{jlk}$. After a complete search, the slice set is $A_{w_{jl}}(u_{ji}) = \{1, 2, 3, 5, 7, 8\}$.

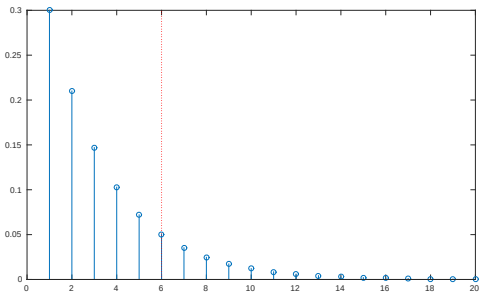

(b) Geometric stick-breaking weights for $N^*_{jl} = 20$. The red dashed line represents the slice variable $N_{ji} = 6$. The slice set is simply $\mathcal{S}_{ji} = \{1, 2, 3, 4, 5, 6\}$.

Figure 1: A visualization of the effect of the u_{ji} snd N_{ji} slice variables are given in Figures 1(a) and 1(b) respectively.

5. Illustrations.

In this section we compare the rPDDP and the PDGSBP models. For the choice of kernel we have a normal model $K(x|\theta) = \mathcal{N}(x|\theta)$ where $\theta = (\mu, \sigma^2)$ and $\tau = \sigma^{-2}$ is the precision. The prior for the means and the precisions for both the rPDDP and PDGSBP models will be independent normals $\mathcal{N}(0, s^{-1})$ and gammas $\mathcal{G}(\tau|\epsilon_1, \epsilon_2)$ respectively i.e.

$$P_0(d\mu, d\tau) = \mathcal{N}(\mu|0, s^{-1})\,\mathcal{G}(\tau|\epsilon_1, \epsilon_2)\,d\mu\,d\tau.$$

Attempting a noninformative prior specification, in all our simulated numerical examples we took $s = \epsilon_1 = \epsilon_2 = 10^{-3}$. A-priori the concentration masses are being set to $c_{jl} \sim \mathcal{G}(a_{jl}, b_{jl})$. To keep the rPDDP and PDGSBP samplers under synchronized prior specifications, for the

geometric probabilities, we set $\lambda_{jl} \sim \mathcal{TG}(a_{jl}, b_{jl})$. In the appendix we show that such prior specifications are valid only for $a_{jl} > 1$. In all our numerical examples, based on real and synthetic data, we took $a_{jl} = b_{jl} = 1.1$. For both models, and for all synthetic data simulations, the hyperparameters of the Dirichlet priors over the selection probabilities (p_{jl}) has been set to $\alpha_{jl} = 1$. Informative prior specifications over the p_{jl}'s are used only in the case of the real data sets. In all cases, the Gibbs samplers run for 11×10^4 iterations leaving the first 10^4 samples as a burn-in period.

First synthetic data example – Equally weighted normal mixtures with unimodal common and idiosyncratic parts: In Figure 1, we show the pairwise dependence of the models for the data for $m = 2, 3, 4$. Each row suggests an equally weighted normal mixture with each circle shaped object indicating the mean of the corresponding normal of unit variance. A vertical line indicates a unimodal dependence between the data sets.

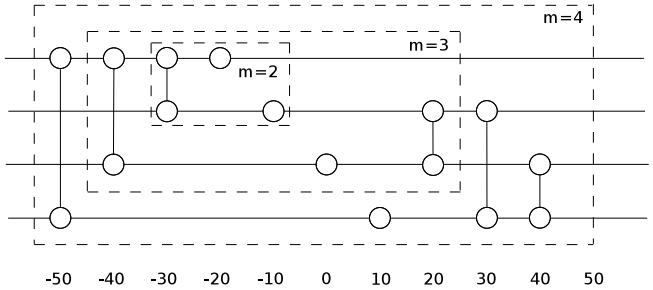

Figure 2: Pairwise dependence and idiosyncratic parts of the synthetic models for the data for $m = 2, 3, 4$.

For example, for $m = 4$ we have:

$$
\begin{aligned}
f_1(x) &\propto \mathcal{N}(x|-20,1) + \mathcal{N}(x|-30,1) + \mathcal{N}(x|-40,1) + \mathcal{N}(x|-50,1)\\
f_2(x) &\propto \mathcal{N}(x|-30,1) + \mathcal{N}(x|-10,1) + \mathcal{N}(x|+20,1) + \mathcal{N}(x|+30,1)\\
f_3(x) &\propto \mathcal{N}(x|-40,1) + \mathcal{N}(x|+20,1) + \mathcal{N}(x|\quad 0,1) + \mathcal{N}(x|+40,1)\\
f_4(x) &\propto \mathcal{N}(x|-50,1) + \mathcal{N}(x|+30,1) + \mathcal{N}(x|+40,1) + \mathcal{N}(x|+10,1).
\end{aligned}
$$

We are taking independently samples of sizes $n_j = 60$ from f_j for $m = 2$, $n_j = 120$ from f_j for $m = 3$, and, $n_j = 200$ from f_j for $m = 4$. In all cases, both models, capture the correct number of modes and they give the correct density estimations.

In Figures 3(a)-(d) we give the histograms of the data sets, for the specific case $m = 4$, which are overladed with the kernel density estimations (KDE's) based on the predictive samples of the f_j's coming from the PDGSBP (solid line) and the rPDDP (dashed line) models. The differences are nearly indistinguishable.

In Table 1 we summarize the mean execution times (MET's) per 10^3 iterations in seconds. The PDGSBP sampler is about three times faster than the rPDDP sampler. The corresponding

MET ratios for $m = 2, 3$ and 4 has been approximated to $2.96, 3.04$ and 3.37 respectively. We can see that the PDGSBP Gibbs sampler gives slightly faster execution times with increasing m. This will become more clear in our next simulated data example, where the average sample size per mode is being kept constant.

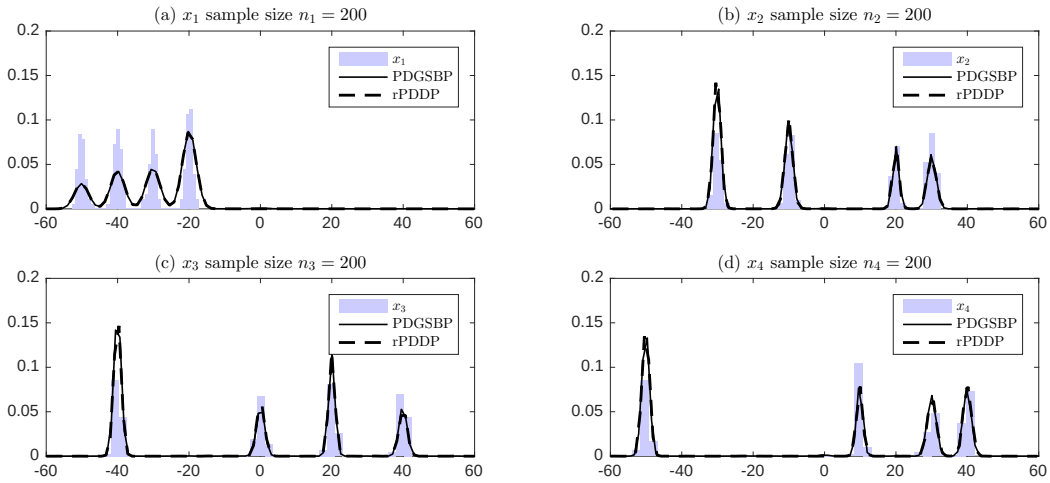

Figure 3: Histograms of data sets coming from equally weighted normal mixtures with unimodal common and idiosyncratic parts for the case $m = 4$. The superimposed KDE's are based on the predictive samples of the PDGSBP and the rPDDP models.

m	Model	Sample size	MET
2	PDGSBP	$n_j = 60$	0.57
	rPDDP		1.68
3	PDGSBP	$n_j = 120$	2.16
	rPDDP		6.57
4	PDGSBP	$n_j = 200$	5.30
	rPDDP		17.87

Table 1: Mean excecution times in seconds per 10^3 iterations.

Second simulated data example – Sparse m-scalable data sets: In this example, to compare the execution times of the PDGSBP and rPDDP models as m increases, we attempt to create m-scalable normal mixture data sets of the lowest possible size. To this respect, we sample independently m data sets from the densities

$$f_j(x) \propto \mathcal{N}(x|(j-1)\xi, 1)\,\mathcal{I}(1 \leq j < m) + \sum_{k=1}^{m-1} \mathcal{N}(x|(k-1)\,\xi, 1)\,\mathcal{I}(j = m),$$

with sample sizes $n_j = n\{\mathcal{I}(1 \leq j < m) + (m-1)\,\mathcal{I}(j = m)\}$. We have set $\xi = 10$ and average sample size per mode $n = 20$, for $m \in \{2, \ldots, 10\}$.

In Table 2 we give the MET's approximations. The faster execution times are due to the sparse nature of the data sets. Now it becomes clear that the PDGSBP sampler becomes faster as m increases. In Figure 4 we depict the average execution times as functions of the dimension m. We can see how fast the two MET-curves diverge as m is increasing. In Figure 5 we give we give the histograms of the data sets, for the specific case $m = 10$, overladed with the KDE's based on the predictive samples of the f_j's coming from the PDGSBP (solid line) and the rPDDP (dashed line) models.

Table 2: MET's based on the sparse m-scalable data sets.

m	MET-rPDDP	MET-PDGSBP	MET-ratio
2	1.10	0.20	5.52
3	3.05	0.53	5.79
4	5.78	1.00	5.78
5	9.18	1.66	5.53
6	13.36	2.42	5.53
7	19.21	3.79	5.15
8	29.20	4.65	6.28
9	43.08	6.27	6.87
10	73.08	9.81	7.45

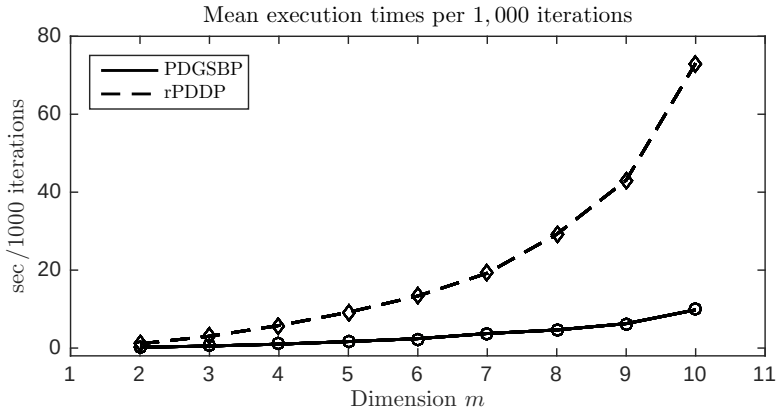

Figure 4: Mean execution times for the two models, based on the sparse m-scalable data sets.

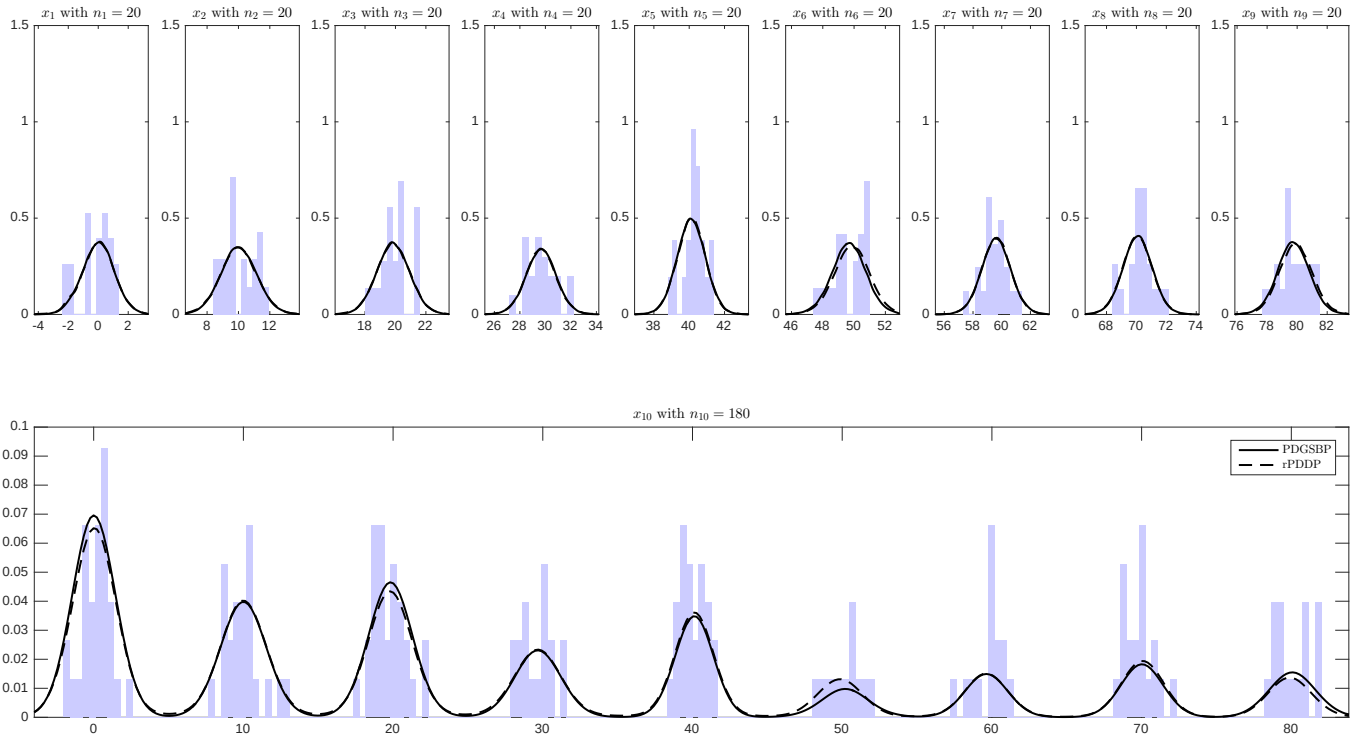
x_1 with $n_1 = 20$
x_2 with $n_2 = 20$
x_3 with $n_3 = 20$
x_4 with $n_4 = 20$
x_5 with $n_5 = 20$
x_6 with $n_6 = 20$
x_7 with $n_7 = 20$
x_8 with $n_8 = 20$
x_9 with $n_9 = 20$
x_{10} with $n_{10} = 180$
PDGSBP
rPDDP

Figure 5: Histograms of sparse m-scalable data sets for the case $m = 10$. The superimposed KDE's are based on the predictive samples of the PDGSBP and the rPDDP models.

Real data example: The data set is to be found at `http://lib.stat.cmu.edu/datasets/pbcseq` and involves data from 310 individuals. We take the observation as SGOT (serum glutamic-oxaloacetic transaminase) level, just prior to liver transplant or death or the last observation recorded, under three conditions on the individual

1. The individual is dead without transplantation.
2. The individual had a transplant.
3. The individual is alive without transplantation.

We normalize the means of all three data sets to zero. Since it is reasonable to assume the densities for the observations are similar for the three categories (especially for the last two), we adopt the models proposed in this paper with $m = 3$. The number of transplanted individuals is small (sample size of 28) so it is reasonable to borrow strength for this density from the other two. In this example, we set the hyperparameters of the Dirichlet priors for the selection probabilities to

$$\alpha_{jl} = \begin{cases} 10, & \text{if } j = l = 1 \text{ or } j = l = 3 \\ 1, & \text{otherwise.} \end{cases}$$

1. In Figure 6(a)-(c) we provide histograms of the real data sets and superimpose the KDE's based on the predictive samples of the PDDP and PDGSBP samplers. The two models give nearly identical density estimations.

2. The estimated a-posteriori selection probabilities are given below

$$\mathbb{E}_{\text{PDGSBP}}(p \,|\, \boldsymbol{x}) = \begin{pmatrix} 0.61 & 0.23 & 0.16 \\ 0.34 & 0.10 & 0.56 \\ 0.08 & 0.12 & 0.80 \end{pmatrix}, \quad \mathbb{E}_{\text{rPDDP}}(p \,|\, \boldsymbol{x}) = \begin{pmatrix} 0.67 & 0.16 & 0.17 \\ 0.29 & 0.15 & 0.56 \\ 0.10 & 0.12 & 0.78 \end{pmatrix}.$$

By comparing the second rows of the selection matrices, we conclude that the the strength of borrowing is slightly larger in the case of PDGSBP model .

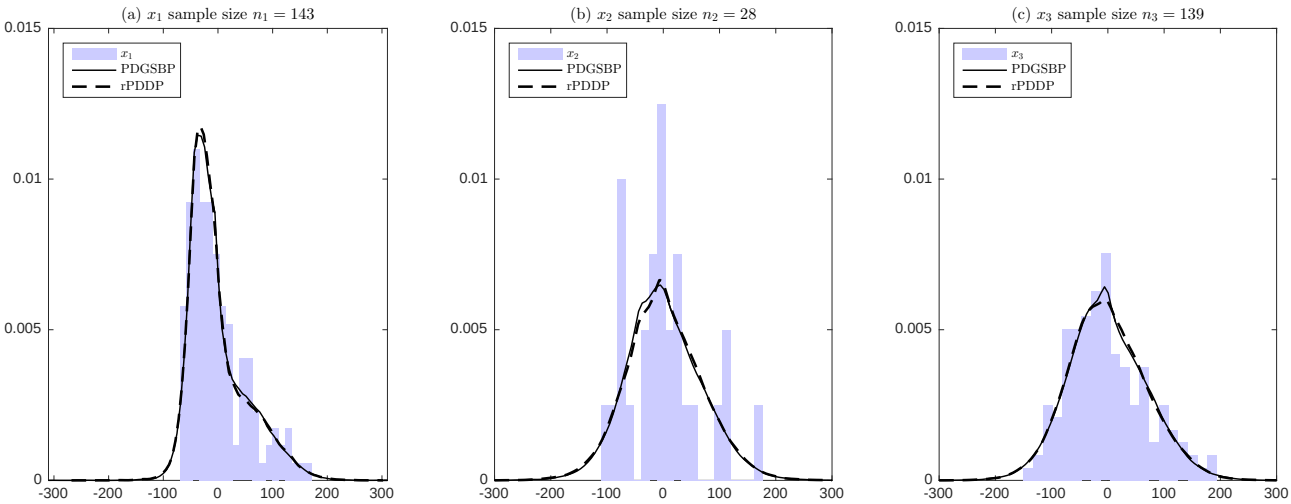

Figure 6: Histograms of the real data sets, and superimposed KDE curves based on the predictive samples of the PDDP and PDGSBP models.

7. Discussion

In this paper we have generalized the GSB process to a multidimensional dependent stochastic process which can be used as a Bayesian nonparametric prior for density estimation in the case of partially exchangeable data sets. The resulting Gibbs sampler is as accurate as its DP based counterpart, faster and by far less complex. The main reason for that, is that the GSB sampled value of the allocation variable d_{ji}, will be an element of the sequential slice set $\mathcal{S}_{ji} = \{1, \ldots, N_{ji}\}$. Thus, there is no need to search the arrays of the weights; we know the state space of the clustering variables in advance. On the other hand, the sampling of d_{ji} in the DP based algorithm will always have one more step; the creation of the slice sets.

For an objective comparison of the execution times of the two models, we have run the two samplers in an a-priori synchronized mode. This, involves the placing of $\mathcal{G}(a_{jl}, b_{jl})$ priors over the DP c_{jl} concentration masses, leading to a more efficient version of the PDDP model introduced in Hatjispyros et al. (2011, 2016).

Appendix

1. Sampling of the concentrations masses for the rPDDP model.

In this case, the random densities (f_j) are represented as finite mixtures of the DP mixtures $g_{jl}(x|\mathbb{P}_{jl})$, where $\mathbb{P}_{jl} \sim \mathcal{DP}(c_{jl}, P_0)$. We randomize the concentrations by letting $c_{jl} \sim \mathcal{G}(a_{jl}, b_{jl})$. Following West (1994) we have the following two specific cases:

A. For $j = l$, the posterior c_{jj}'s will be affected only by the size of the data set $\boldsymbol{x}_j$ and the number of unique clusters for which $\delta_{ji} = \mathbf{e}_j$. Letting

$$\rho_{jj} = \#\{d_{jj} : \delta_{ji} = \mathbf{e}_j, 1 \leq i \leq n_j\},$$

we have

$$\begin{aligned}
&\beta \sim \mathcal{B}e(c_{jj} + 1, n_j) \\
&c_{jj} \,|\, \beta, \rho_{jj} \;\sim\; \pi_\beta \,\mathcal{G}(a_{jj} + \rho_{jj}, b_{jj} - \log\beta) + (1 - \pi_\beta)\,\mathcal{G}(a_{jj} + \rho_{jj} - 1, b_{jj} - \log\beta)
\end{aligned}$$

with the weights π_β satisfying $\frac{\pi_\beta}{1-\pi_\beta} = \frac{a_{jj}+\rho_{jj}-1}{n_j(b_{jj}-\log\beta)}$.

B. For $j \neq l$, the posterior c_{jl}'s will be affected by the size of the data sets $\boldsymbol{x}_j$ and $\boldsymbol{x}_l$ and the cumulative number of unique clusters d_{ji} for which $\delta_{ji} = \mathbf{e}_l$ and the unique clusters d_{li} for which $\delta_{li} = \mathbf{e}_j$. Letting

$$\rho_{jl} = \#\{d_{ji} : \delta_{ji} = \mathbf{e}_l, 1 \leq i \leq n_j\} + \#\{d_{li} : \delta_{li} = \mathbf{e}_j, 1 \leq i \leq n_l\},$$

it is that

$$\begin{aligned}
&\beta \sim \mathcal{B}e(c_{jl} + 1, n_j + n_l) \\
&c_{jl} \,|\, \beta, \rho_{jl} \;\sim\; \pi_\beta \,\mathcal{G}(a_{jl} + \rho_{jl}, b_{jl} - \log\beta) + (1 - \pi_\beta)\,\mathcal{G}(a_{jl} + \rho_{jl} - 1, b_{jl} - \log\beta),
\end{aligned}$$

with the weights π_β satisfying $\frac{\pi_\beta}{1-\pi_\beta} = \frac{a_{jl}+\rho_{jl}-1}{(n_j+n_l)(b_{jl}-\log\beta)}$.

Bear in mind that $\rho_{jl} = 0$ is always a possibility, so that we impose $a_{jl} > 1$.

2. Sampling of the geometric probabilities for the PDGSBP model.

In this section we provide the full conditionals for the geometric probabilities λ_{jl} under beta conjugate and transformed gamma nonconjugate priors. We let

$$S_{jl} = \sum_{i=1}^{n_j} \mathcal{I}(\delta_{ji} = \mathbf{e}_l) \quad \text{and} \quad S'_{jl} = \sum_{i=1}^{n_j} \mathcal{I}(\delta_{ji} = \mathbf{e}_l)(N_{ji} - 1).$$

A. For the choice of prior $\lambda_{jl} \sim \mathcal{B}e(a_{jl}, b_{jl})$, for $l = j$ it is that

$$f(\lambda_{jj}|\cdots) = \mathcal{B}e(\lambda_{jl}|a_{jj} + 2S_{jj}, b_{jj} + S'_{jj}),$$

also, for $l \neq j$ we have

$$f(\lambda_{jl}|\cdots) = \mathcal{B}e(\lambda_{jl}|a_{jl} + 2(S_{jl} + S_{lj}), b_{jl} + S'_{jl} + S'_{lj}).$$

B. For the choice of prior $\lambda_{jl} \sim \mathcal{TG}(a_{jl}, b_{jl})$, for $l = j$ it is that

$$f(\lambda_{jj}|\ldots) \propto \lambda_{jj}^{2S_{jj}-a_{jj}-1}(1-\lambda_{jj})^{S'_{jj}+a_{jj}-1}e^{-b_{jj}/\lambda_{jj}}\,\mathcal{I}(0<\lambda_{jj}<1).$$

To sample from this density, we include the positive auxiliary random variables ν_1 and ν_2 such that

$$f(\lambda_{jj},\nu_1,\nu_2|\cdots) \propto \lambda_{jj}^{2S_{jj}-a_{jj}-1}\mathcal{I}\left(\nu_1 < (1-\lambda_{jj})^{S'_{jj}+a_{jj}-1}\right)\mathcal{I}\left(\nu_2 < e^{-b_{jj}/\lambda_{jj}}\right)\mathcal{I}(0<\lambda_{jj}<1).$$

The full conditionals for ν_1, ν_2 are uniforms

$$f(\nu_1|\cdots) = \mathcal{U}\left(\nu_1|0,(1-\lambda_{jj})^{S'_{jj}+a_{jj}-1}\right) \quad \text{and} \quad f(\nu_2|\cdots) = \mathcal{U}\left(\nu_2|0,e^{-b_{jj}/\lambda_{jj}}\right),$$

and the new full conditional for λ_{jj} becomes

$$f(\lambda_{jj}|\nu_1,\nu_2,\ldots) \propto \lambda_{jj}^{2S_{jj}-a_{jj}-1}\begin{cases}\mathcal{I}\left(-\frac{b_{jj}}{\log\nu_2} < \lambda_{jj} < 1-\nu_1^{1/L_{jj}}\right) & L_{jj}\geq 0\\ \mathcal{I}\left(\max\left\{-\frac{b_{jj}}{\log\nu_2}, 1-\nu_1^{1/L_{jj}}\right\} < \lambda_{jj} < 1\right) & L_{jj}<0,\end{cases}$$

where we have set $L_{jj} = S'_{jj} + a_{jj} - 1$. We can sample from this density using the inverse cummulative distribution function technique. Also, for $l \neq j$ we apply the same embedded Gibbs sampling technique to the full conditional density

$$f(\lambda_{jl}|\cdots) \propto \lambda_{jl}^{2(S_{jl}+S_{lj})-a_{jl}-1}(1-\lambda_{jl})^{S'_{jl}+S'_{lj}+a_{jl}-1}e^{-b_{jl}/\lambda_{jl}}\,\mathcal{I}(0<\lambda_{jl}<1).$$

References.

Bulla, P., Muliere, P. and Walker, S.G. (2009). A Bayesian nonparametric estimator of a multivariate survival function. *Journal of Statistical Planning and Inference* **139**, 3639–3648.

De Iorio, M., Müller, P., Rosner, G.L. and MacEachern, S.N. (2004). An ANOVA model for dependent random measures. *Journal of the American Statistical Association* **99**, 205–215.

Dunson, D.B. and Park, J.H. (2008). Kernel stick–breaking processes. *Biometrika* **95**, 307–323.

Ferguson, T.S. (1973). A Bayesian analysis of some nonparametric problems. *Annals of Statistics* **1**, 209–230.

Foti, N.J. and Williamson, S. (2014). A survey of non–exchangeable priors for Bayesian nonparametric models. *Pattern Analysis and Machine Intelligence,* IEEE Transactions **37**(2), 359–371.

Fuentes–Garcia, R., Mena, R.H., Walker, S.G. (2009). A nonparametric dependent process for Bayesian regression *Statistics and Probability Letters* **79**, 1112–1119.

Fuentes–Garcia, R., Mena, R.H., Walker, S.G. (2010). A new Bayesian nonparametric mixture model. *Comm.Statist.Simul.Comput* **39**, 669–682.

Griffin, J.E. and Steel, M.F.J. (2006). Order–based dependent Dirichlet processes. *Journal of the American Statistical Association* **101**, 179–194.

Griffin, J.E., Kolossiatis, M. and Steel, M.F.J. (2013). Comparing distributions by using dependent normalized ranom–measure mixtures. *Journal of the Royal Statistical Society, Series B* **75**, 499–529.

Hatjispyros, S.J., Nicoleris, T. and Walker, S.G. (2011). Dependent mixtures of Dirichlet processes. *Computational Statistics and Data Analysis* **55**, 2011–2025.

Hatjispyros, S.J., Nicoleris, T. and Walker, S.G. (2016). Dependent random density functions with common atoms and pairwise dependence. *Computational Statistics and Data Analysis* **101**, 236–249.

Hatjispyros, S.J., Merkatas, C. (2015). Bayesian Nonparametric Reconstruction and Prediction of Nonlinear Dynamic Systems with Geometric Stick Breaking Noise *submitted,* arXiv:1511.00154.

Lijoi, A., Nipoti, B. and Prüenster, I. (2014). Dependent mixture models: clustering and borrowing information. *Computational Statistics and Data Analysis* **71**, 17–433.

Kolossiatis, M., Griffin, J.E. and Steel, M.F.J. (2013). On Bayesian nonparametric modelling of two correlated distributions. *Statistics and Computing* **23**, 1–15.

Lo, A.Y. (1984). On a class of Bayesian nonparametric estimates I. Density estimates. *Annals of Statistics* **12**, 351–357.

MacEachern, S.N. (1999). Dependent nonparametric processes. In *"Proceedings of the Section on Bayesian Statistical Science"* pp. 50-55. American Statistical Association.

Mena, R.H., Ruggiero, M. and Walker, S.G. (2011). Geometric stick–breaking processes for continuous–time Bayesian nonparametric modeling. *Journal of Statistical Planning and Inference* **141** (9), 3217–3230.

Mena, R.H., Walker, S.G. (2015). On the Bayesian mixture model and identifiability. *Journal of Computational and Graphical Statistics* **24** (4), 1155-1169.

Muliere, P. and Tardella, L. (1998). Approximating distributions of random functionals of Ferguson–Dirichlet priors. *Can J Statistics* **29** 283–297, doi: 10.2307/3315511

Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica* **4** 639–650

Teh, Y.,W., Jordan, M.,I., Beal, M.,J., Blei, D.,M., (2006). Hierarchical Dirirchlet Processes. *Journal of the American Statistical Association* 101 **476** 1566–1581

Walker, S.G. (2007). Sampling the Dirichlet mixture model with slices *Communications in Statistics* **36** 45–54

West, M. (1992). Hyperparameter estimation in Dirichlet process mixture models. *Technical report* **92-A03**, Duke University, ISDS